From symmetric fundamental expansions
to Schur positivity

Austin Roberts

Department of Mathematics

Highline College

Des Moines, WA 98198, USA

August 31, 2015

Abstract

We consider families of quasisymmetric functions with the property that if a symmetric function f is a positive sum of functions in one of these families, then f is necessarily a positive sum of Schur functions. Furthermore, in each of the families studied, we give a combinatorial description of the Schur coefficients of f. We organize six such families into a poset, where functions in higher families in the poset are always positive integer sums of functions in each of the lower families. This poset includes the Schur functions, the quasisymmetric Schur functions, the fundamental quasisymmetric generating functions of shifted dual equivalence classes, as well as three new families of functions — one of which is conjectured to be a basis of the vector space of quasisymmetric functions. Each of the six families is realized as the fundamental quasisymmetric generating functions over the classes of some refinement of dual Knuth equivalence. Thus, we also produce a poset of refinements of dual Knuth equivalence. In doing so, we define quasi-dual equivalence to provide classes that generate quasisymmetric Schur functions.

1

1 Introduction

The problem of how to express a symmetric function in terms of the basis of Schur functions arises prominently in many fields, including algebraic combinatorics, representation theory, and statistical mechanics, amongst others. For instance, showing that a function is a positive integer sum of Schur functions (Schur positive) is equivalent to showing that the function corresponds to a representation of the general linear group, where the coefficients of said sum give the multiplicities of irreducible sub representations. See [Sagan, 2001] or [Stanley, 1999]) for a treatment. In many cases, such as Macdonald polynomials or plethysms of Schur functions, a symmetric function has a known expansion in terms of the fundamental quasisymmetric functions while an explicit expansion over the Schur functions remains elusive (see [Haglund et al., 2005] and [Loehr and Warrington, 2012]).

In this paper, we consider families of quasisymmetric functions with the property that if a symmetric function f is a positive sum of functions in one of these families, then f is necessarily a positive sum of Schur functions. Furthermore, in each of the families studied, we give a combinatorial description of the Schur coefficients of f. We organize six such families into a poset, where functions in higher families in the poset are positive integer sums of functions in each of the lower families. This poset includes the Schur functions, quasisymmetric Schur functions, the fundamental quasisymmetric generating functions of shifted dual equivalence classes $\{f^{(h)}\}$, as well as three new families of functions $\{f^{(k)}\}$ for $k = 0, 1, 2$ — conjecturing that $\{f^{(2)}\}$ forms a basis for the vector space of quasisymmetric functions.

The poset of functions is first realized using a poset of equivalence relations. Each of the families is defined as sums of fundamental quasisymmetric functions over equivalence classes on standard Young tableaux of fixed partition shape λ, SYT(λ). A higher position in our poset represents a courser relation. We then use the Robinson-Schensted-Knuth (RSK) correspondence to turn equivalence classes on tableaux into equivalence classes on permutations, realizing each equivalence relation as some restriction of dual Knuth equivalence. An illustration of said poset of equivalence classes, the related poset of quasisymmetric functions, and the generators of these classes can be found in Figure 1.

While detailed definitions will be given later, we can describe the main results of the paper as follows. In [Gessel, 1984], Gessel used the set SYT(λ) to express the Schur function s_λ as,

$$s_\lambda = \sum_{T \in \text{SYT}(\lambda)} F_{\text{ID}(T)}(X), \tag{1.1}$$

where $F_{\text{ID}(T)}(X)$ is the fundamental quasisymmetric function

$$F_{\text{ID}(T)}(X) = \sum_{\substack{j_1 \leq \cdots \leq j_n \\ j_k = j_{k+1} \implies j \notin \text{ID}(T)}} x_{j_1} \cdots x_{j_n}. \tag{1.2}$$

Here, ID(T) is the inverse descent set of T. We define three equivalence relations $\equiv^0$, $\equiv^1$, and $\equiv^2$ on SYT(λ), each a refinement of the next. One particularly nice equivalence class of $\equiv^0$
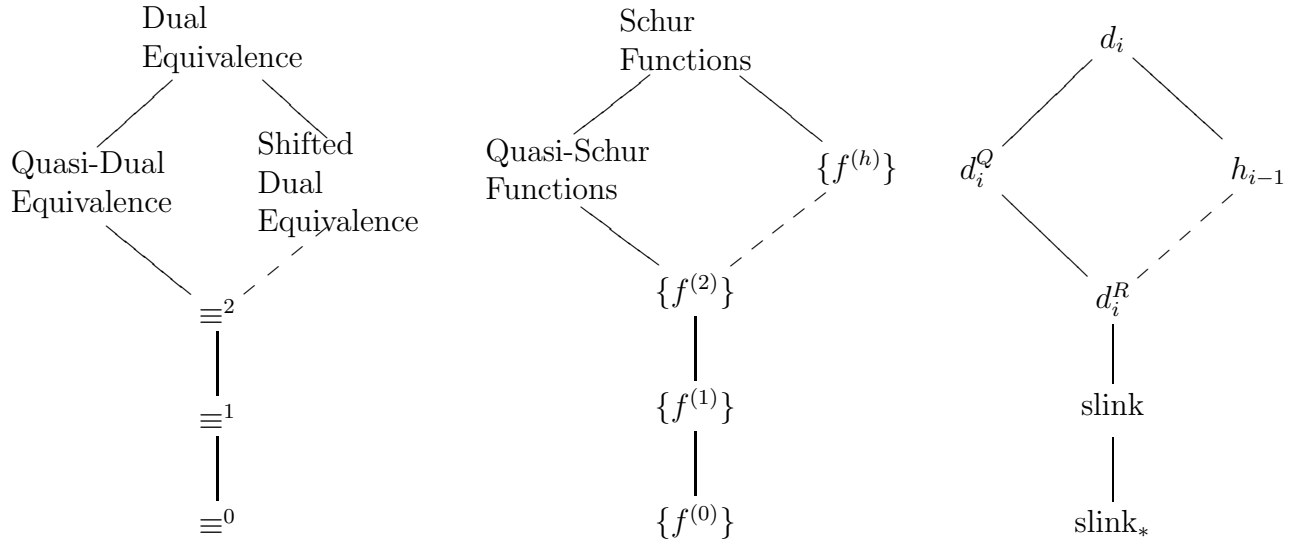
2

Figure 1: At left, the poset of equivalence classes. At center, the associated quasisymmetric functions of the equivalence class. At right, the generators of the equivalence relations. The dashed lines denotes that a transformation is required, reversing all permutations in each equivalence class for the left and right posets.

and $\equiv^1$ is the class of a single element, the *superstandard tableau* U_λ. Here, λ is a partition of n and $U_\lambda \in \mathrm{SYT}(n)$ is the standard Young tableau attained by filling each row of λ, in order, with as small of values as possible. The main result can then be stated as follows.

Theorem 1.1. *For $k = 0, 1, 2$, let $\mathcal{C} \subset \mathrm{SYT}(n)$ be the disjoint union of equivalence classes of $\equiv^k$ such that $f = \sum_{T \in \mathcal{C}} F_{\mathrm{ID}(T)}$ is a symmetric function. Then*

$$f = \sum_{\lambda \vdash n} c_\lambda s_\lambda,$$

where c_λ is the multiplicity of U_λ in $\mathcal{C}$.

Corollary 2.18 then generalizes this theorem to equivalence classes of permutations, where each relation becomes a refinement of dual Knuth equivalence. All of these relations on permutations has the added benefit of commuting with Knuth equivalences, (or, equivalently, jeu de taquin), as stated in Proposition 2.16. In Conjecture 2.20, we further propose that the family of fundamental quasisymmetric generating functions over $\equiv^2$, $\{f^{(2)}\}$, form a basis for the quasisymmetric functions.

We then turn our attention to two applications. First is the set of quasisymmetric functions, which — as the name suggests — are a quasisymmetric analogue of the Schur functions. They were introduced in [Haglund et al., 2011] and have since been studied for there relation to Demazure atoms, the Littlewood-Richardson rule, the RSK correspondence, and 0-Hecke algebras, amongst others (see [Luoto et al., 2013] for an overview of the topic). Quasisymmetric

3

Schur functions can be realized as a sum over standard reverse composition tableaux of a fixed composition shape, SRCT(α). In Proposition 3.3, we describe a transitive action on SRCT(α). Using a result in [Mason, 2008], SRCT(α) can be mapped to a subset of the standard reverse tableaux of a fixed partition shape SRT(λ). We show in Corollary 3.5 that Mason's bijection sends our transitive action on SRCT(α) to a subset of dual Knuth equivalences, which we term *quasi-dual equivalences*. In Lemma 3.7, we show that the quasi-dual equivalence is a coarsening of $\equiv^k$, for each $k = 0, 1, 2$. Hence the quasisymmetric Schur functions are positive integer sums of functions in $\{f^{(k)}\}$, as stated in Theorem 3.7.

Finally, we consider shifted dual equivalence. Our presentation is most closely related to [Haiman, 1992], though it was was originally studied simultaneously in [Sagan, 1987] and [Worley, 1984]. Shifted dual equivalence is related to enumerative properties of reduced words in Lie type B, Stanley symmetric functions, the P and Q-Schur functions, and shifted dual equivalence graphs (also see [Billey et al., 2014], [Stanley, 1984], and [Stanley, 1999, Ch. 7]). Similar to with the quasisymmetric Schur functions, Proposition 4.4 states that the generators for shifted dual equivalence strictly contain the generators for $\equiv^2$ after applying a simple involution. Hence, the fundamental quasisymmetric generating functions over shifted dual equivalence classes may be added to our poset, as stated in Theorem 4.5. In Proposition 4.8, we further show that the set of row reading words of shifted standard Young tableaux with a fixed shape comprise an equivalence class of $\equiv^2$.

The paper is organized as follows. After preliminary lemmas and definitions, Section 2 defines the equivalence relations $\equiv^k$ and proves Theorem 1.1. Section 2.3 generalizes results to permutations, with the main result generalized in Corollary 2.3. Section 3 is dedicated to the previously mentioned results related to quasisymmetric Schur functions. Finally, Section 4, which is independent of Section 3, is dedicated to proving the results related to shifted dual equivalence.

2 The equivalence relations

2.1 Preliminaries

A composition $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$ is a finite sequence of positive integers. We write $\alpha \vDash n$ and say that α is a *composition of n* if $\sum_i \alpha_i = n$. A partition $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ is a weakly decreasing composition. We write $\lambda \vdash n$ and say λ is a partition of n if $\sum_i \lambda_i = n$. In this paper, λ will always be a partition of n, α will always be a composition of n, and $\lambda(\alpha)$ is the composition achieved by sorting the parts of α into weakly decreasing order.

The *diagram* of a composition or a partition will always be given in French notation. That is, the diagram is given by a set of left justified *cells*, drawn as boxes, in the Cartesian plane, where the i^{th} row from bottom to top has α_i or λ_i cells, respectively. The bottom left cell can be assumed to be the origin. The underlying composition or partition is then referred to as the *shape* of the diagram. Given a partition λ, the *conjugate* of λ, denoted λ', is the partition whose i^{th} column has λ_i many cells.

A *filling T* of a diagram is a function that assigns a positive integer to each box of the diagram. A filling is *standard* if it uses every number in $[n] = \{1, \ldots, n\}$ exactly once. The set

of *standard Young tableaux* of shape λ, denoted SYT(λ), is the set of standard fillings of λ that are increasing up columns and across rows from left to right. The union of SYT(λ) over all $\lambda \vdash n$ is denoted SYT(n). The *superstandard tableau* of shape λ, denoted U_λ, is the standard Young tableau formed by placing the numbers 1 through n, in order, in the lowest row possible. The *row reading word* of a filling, rw(T), is obtained by reading the values of each row from left to right, starting with the top row and continuing down. The row reading word will always be used for standard Young tableaux, as well as for ordering their cells, but other reading words will be introduced for other types of fillings as needed.

In this paper, permutations are always given in one-line notation. Given a permutation $\pi \in S_n$, the *inverse descent set* of π, ID(π), is the set of $i \in [n-1]$ such that i occurs after $i+1$ in π. Alternately, we can encode ID(π) with the composition $\beta(\pi) = (\beta_1(\pi), \ldots, \beta_k(\pi)) \vDash n$, letting β_i be the difference between the i and $i-1^{th}$ numbers in ID(π) $\cup \{n\}$, where the zeroth number is always treated as 0. Given a standard filling T, define ID(T) and $\beta(T)$ via the reading word of T. Define the i^{th} *run of* T, $T|_{(a,b]}$, to be the restriction of T to the set of cells with values in the integer interval $(a, b]$, where a and b are the $i-1$ and i^{th} numbers in ID(T) $\cup \{n\}$, respectively. Notice that $b - a = \beta_i$. Similarly, the *first j runs of* T is the standard Young tableau achieved by taking the union of the first j runs of T. Notice that for any $T \in$ SYT(λ), the cells of the runs of T fully determine T. For this reason, it can be helpful to consider the *unstandard Young tableau*, UYT, achieved by placing i's in each cell of the i^{th} run. See Figure 2 for examples.

Figure 2: Two tableaux in SYT(λ) with their associated fillings in UYT(λ) to their right. Here, $\lambda = (4, 4, 1)$. At left, U_λ with rw(U_λ) = 956781234, ID(U_λ) = $\{4, 8\}$, and $\beta(T) = (4, 4, 1)$. At right, a tableau T with rw(T) = 634891257, ID(T) = $\{2, 5, 7\}$ and $\beta(T) = (2, 3, 2, 2)$.

The RSK correspondence provides a bijection that sends each $\pi \in S_n$ to an ordered pair of standard Young tableaux $(P(\pi), Q(\pi))$, where $P(\pi)$ and $Q(\pi)$ are both in SYT(λ) for some $\lambda \vdash n$. Here, $P(\pi)$ is termed the *insertion* tableau and $Q(\pi)$ is termed the *recording tableau*. Two permutations with the same P tableau are *Knuth equivalent*, while two permutations with the same Q tableau are *dual Knuth equivalent* or just *dual equivalent*. The map from π to $(P(\pi), Q(\pi))$ can be achieved via the 'row bumping algorithm' or 'jeu de taquin' and is more generally defined on words, but we will restrict our attention to permutations. We will assume familiarity with the RSK correspondence, though a treatment can be found in [Sagan, 2001].

Let ϕ be a function on standard Young tableaux that restricts to an involution on each SYT(λ). Define the action of ϕ on S_n *via the insertion tableau* to be the function that sends $\pi \in S_n$ to the unique π' such that $P(\pi') = \phi \circ P(\pi)$ and $Q(\pi') = Q(\pi)$.

In the context of permutations, the role of U_λ will be replaced by the set of *standardized Yamanouchi words* of shape λ, defined as

$$\text{SYam}(\lambda) = \{\pi \in S_n \colon P(\pi) = U_\lambda\} \tag{2.1}$$

5

As an aside, we may also algorithmically generate SYam(λ). First generate the set of words with λ_i many i's such that when reading in reverse order there are always weakly more i's than $i+1$'s. These are the Yamanouchi words of weight λ. To achieve SYam(λ), turn each word into a permutation with the same relative order, where an i that occurs earlier in reading order is considered smaller than an i with a larger index.

We may generate *dual equivalence classes* of permutations as follows. Given a permutation in S_n expressed in one-line notation, define an *elementary dual equivalence* as an involution d_i that interchanges the values $i-1, i$, and $i+1$ as

$$
\begin{aligned}
d_i(\ldots i \ldots i-1 \ldots i+1 \ldots) &= (\ldots i+1 \ldots i-1 \ldots i \ldots), \\
d_i(\ldots i \ldots i+1 \ldots i-1 \ldots) &= (\ldots i-1 \ldots i+1 \ldots i \ldots),
\end{aligned}
\tag{2.2}
$$

and as the identity when i is between $i-1$ and $i+1$. Two words are then dual equivalent if one may be transformed into the other by successive elementary dual equivalences. As an example, 21345 is dual equivalent to 51234 because $d_4(d_3(d_2(21345))) = d_4(d_3(31245)) = d_4(41235) = 51234$.

For a permutation $\pi \in S_n$, let $\pi|_I$ be the subword consisting of values in the interval I. Let $\mathrm{fl}(\pi|_I) \in S_{|I|}$ be the permutation with the same relative order as $\pi|_I$. Here fl is the *flattening operator*. For example $\mathrm{fl}(31245|_{[2,4]}) = \mathrm{fl}(324) = 213$. By considering the action of d_i on $\mathrm{fl}(\pi|_{[i-1,i+1]})$, we may express all nontrivial actions as

$$
x1y \qquad x3y,
\tag{2.3}
$$

where d_i swaps the values x and y in $\mathrm{fl}(\pi|_{[i-1,i+1]})$ while fixing all values of π not in $[i-1, i+1]$.

Elementary dual equivalences act on standard Young tableaux via their row reading words, as shown in Figure 3. The next theorem expresses the transitivity of this action.
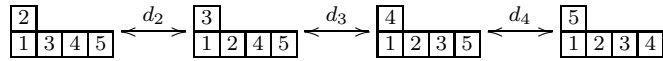

Figure 3: The action of d_i on SYT$((4,1))$.

Theorem 2.1 ([Haiman, 1992, Prop. 2.4]). *Two standard Young tableaux on partition shapes are dual equivalent if and only if they have the same shape.*

In fact, d_i is defined so that

$$
P \circ d_i(\pi) = d_i \circ P(\pi) \quad \text{and} \quad Q \circ d_i(\pi) = Q(\pi).
\tag{2.4}
$$

Furthermore

$$
\mathrm{ID}(P(\pi)) = \mathrm{ID}(\pi).
\tag{2.5}
$$

The *Knuth equivalences* act similarly on permutations, with the roles of P and Q reversed. Define the generator

$$
K_i(\pi) = (d_i(\pi^{-1}))^{-1}.
\tag{2.6}
$$

6

More explicitly, to perform K_i, consider $\pi_{[i-1,i+1]}$ written using the values $x < y < z$. Then K_i permutes the values in $\pi_{[i-1,i+1]}$ by taking yxz to yzx; xzy to zxy; and fixing xyz as well as zyx. For example, $K_4(42153) = 42513$. These involutions interact with the RSK correspondence by

$$P \circ K_i(\pi) = P(\pi) \quad \text{and} \quad Q \circ K_i(\pi) = d_i \circ Q(\pi). \tag{2.7}$$

The previous facts about dual equivalence were shown in [Haiman, 1992], while a more general treatment may be found in [Sagan, 2001].

We will be using tableaux as an enumerative tool for symmetric function calculations. Traditionally, the symmetric functions are defined as the fixed set of formal power series in variables $x_1, x_2, \ldots$ under permutations of their indices. For our purposes, it is enough to define the symmetric functions of degree n as the real vector space generated by the basis of Schur functions s_λ, where $\lambda \vdash n$. We will additionally care about the spanning set of *composition Schur functions*, s_α, where $\alpha \vDash n$. Each s_α is equal to 0, s_λ, or $-s_\lambda$ for a unique $\lambda \vdash n$ attained from α. To be more specific, if β is the result of taking α and replacing some (α_i, α_{i+1}) with $(\alpha_{i+1} - 1, \alpha_i + 1)$, then $s_\alpha = -s_\beta$. In particular, if there exists an i such that $\alpha_i + 1 = \alpha_{i+1}$, then $s_\alpha = 0$. If $s_\alpha \neq 0$, a series of such swaps of α_i and α_{i+1} may always be used to transform α into a unique partition λ.

The relationship between composition Schur functions and (partition) Schur functions is sometimes called the *slinky correspondence* for the following reason. If α is represented by a diagram and $s_\alpha \neq 0$, then we may relate s_α to a unique s_λ, where $\lambda \vdash n$ by letting gravity pull each row down to create a partition shape, as in Figure 4. If this process does not yield a partition, the result is 0. Otherwise, the partition shape is λ, and the sign in front of the Schur function s_λ is $(-1)^d$, where d is the sum of the number or rows each part of the composition was shifted down to make the partition. The term *slinky* is meant to evoke how each row *slinks* downward.
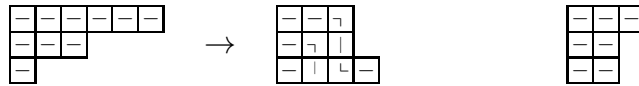

Figure 4: The slinky correspondence showing $s_{(1,3,6)} = (-1)^3 s_{(4,3,3)} = -s_{(4,3,3)}$ and $s_{(2,2,3)} = 0$.

The following relationship between fundamental quasisymmetric functions and composition Schur functions will prove crucial for our purposes.

Lemma 2.2 ([Egge et al., 2010]). *If f is a symmetric function such that*

$$f = \sum_{\alpha \vDash n} c_\alpha F_\alpha, \quad \text{then} \quad f = \sum_{\alpha \vDash n} c_\alpha s_\alpha, \tag{2.8}$$

where c_α are constants.

2.2 The equivalence relations on tableaux

Definition 2.3. Consider $T \in \mathrm{SYT}(\lambda)$. Define $\mathrm{slink}(T)$ and $\mathrm{slink}_*(T)$ as follows.

- Let j be the minimal number such that the first j runs of T do not form a superstandard tableau. If such a j does not exist ($T = U_\lambda$), then slink and slink$_*$ act as the identity.

- Construct slink$(T) \in \mathrm{SYT}(\lambda)$ by permuting the cells of the $j-1$ and j^{th} runs that occur below the j^{th} row, giving the first $\beta_j(T) - 1$ of these cells to the $j - 1^{th}$ run. In particular,

$$\beta_{j-1,j}(\mathrm{slink}(T)) = (\beta_j(T) - 1, \beta_{j-1}(T) + 1). \tag{2.9}$$

- Let μ be the shape of the first j runs of T, and let i be the minimal number such that

$$\mu_{i+1}(T) \leq \beta_j(T) + i - j. \tag{2.10}$$

- Construct slink$_*(T) \in \mathrm{SYT}(\lambda)$ by permuting the cells of the i^{th} and j^{th} runs that occur below the j^{th} row, giving the first $\beta_j(T) + i - j$ of these cells the i^{th} run. In particular,

$$\beta_{i,j}(\mathrm{slink}_*(T)) = (\beta_j(T) + i - j, \beta_i(T) + j - i). \tag{2.11}$$

Notice that the definition above refers to permuting cells in runs, dictating what the values in those cells will have to be. See Figure 5 for examples. A few ramifications of this definition should be pointed out.

Lemma 2.4. *For any partition λ and $T \in \mathrm{SYT}(\lambda)$, the following hold.*

1. *slink(T) and slink$_*(T)$ are well defined and in $\mathrm{SYT}(\lambda)$.*

2. *If $T \neq U_\lambda$, then $s_{\beta(\mathrm{slink}(T))} = -s_{\beta(T)}$.*

3. *If $i = j - 1$ in Definition 2.3, then slink$(T) = $ slink$_*(T)$*

Proof. The proofs are direct consequences of the definition and are recommended for the reader. $\square$

We may now define two equivalence classes that will motivate much of the rest of this paper.

Definition 2.5. Define $\equiv^0$ and $\equiv^1$ as the equivalence relations on standard Young tableaux generated by slink$_*$ and slink, respectively.

Examples of equivalence classes can be found in Figure 5. The relationships between slink$_*$ and $\equiv^1$ are further illuminated by the following lemmas.

Lemma 2.6. *For any partition λ and $T \in \mathrm{SYT}(\lambda)$ such that $T \neq U_\lambda$, let i and j be as in Definition 2.3.*

1. *If $i = j - 1$, then slink fixes i and j,*

2. *If $i < j - 1$, then slink fixes i and lowers j by 1,*
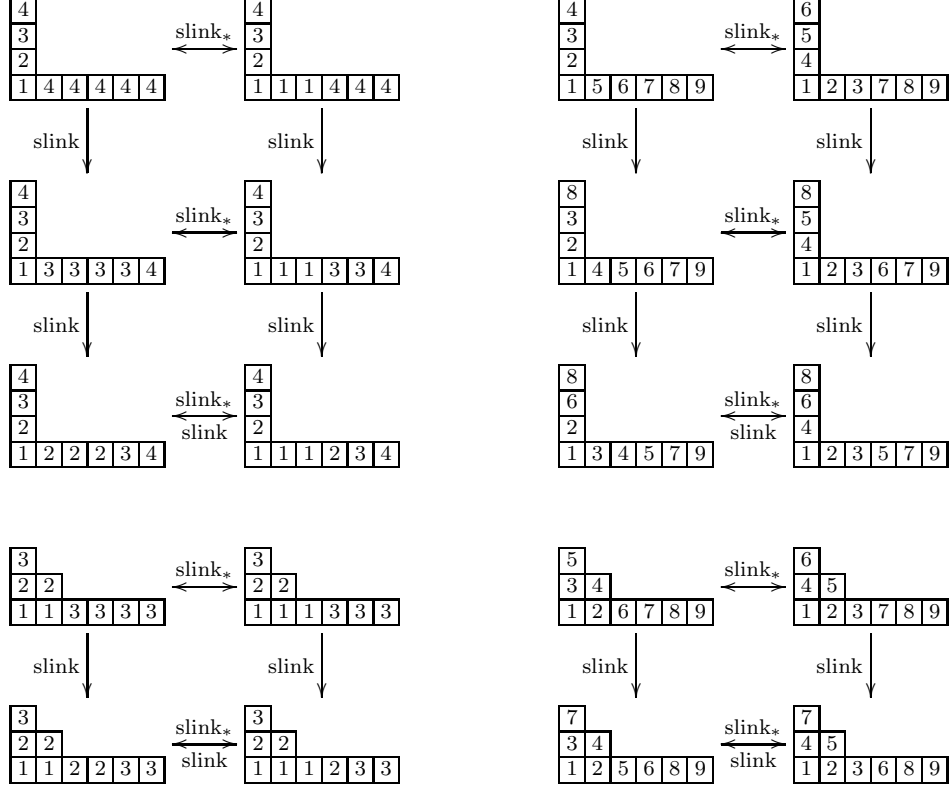
3. *slink$_*$ fixes i and j.*

Figure 5: Two equivalence classes of $\equiv^1$ using $\mathrm{SYT}(\lambda)$ on the right and their representation in $\mathrm{UYT}(\lambda)$ on the left. The $\equiv^0$ classes are the vertex sets of edges labeled slink_*.

Proof. Let T, i and j be as in the statement of the Lemma. Let μ be as in Definition 2.3, where we use $\mu(T), \mu(\mathrm{slink}(T))$, and $\mu(\mathrm{slink}_*(T))$ to differentiate the possibly different shapes associated to each tableau. Consider the three parts of the lemma in order.

Part 1: Assume $i = j - 1$. By Part 3 of Lemma 2.4, $\mathrm{slink}(T) = \mathrm{slink}_*(T)$. The result is thus a special case of Part 3, which is proven below.

Part 2: Assume $i < j - 1$. By (2.10) and (2.9),

$$\mu_i(T) \leq \beta_j(T) + i - j < \beta_j(T) - 1 = \beta_{j-1}(\mathrm{slink}(T)). \tag{2.12}$$

Thus, the $j - 1^{th}$ run of $\mathrm{slink}(T)$ must take all of the cells of the j^{th} run of T that occur weakly between the i and $j - 1^{th}$ rows, as well as at least one cell below the i^{th} row. Hence, the first $j - 1$ runs of $\mathrm{slink}(T)$ do not form a superstandard tableau, and the j value must be lowered by 1. Further, $\mu_{i,i+1}(T) = \mu_{i,i+1}(\mathrm{slink}(T))$.

To show that slink fixes i, it suffices to use (2.9) to inspect (2.11) for both i and $i - 1$,

$$\begin{aligned}
\mu_{i+1}(\mathrm{slink}(T)) &= \mu_{i+1}(T) \leq \beta_j(T) + i - j \\
&= \beta_{j-1}(\mathrm{slink}(T)) + i - (j - 1),
\end{aligned} \tag{2.13}$$

9

$$\begin{aligned}
\mu_i(\mathrm{slink}(T)) =& \mu_i(T) > \beta_j(T) + (i-1) - j \\
=& \beta_{j-1}(\mathrm{slink}(T)) + (i-1) - (j-1).
\end{aligned} \tag{2.14}$$

Thus, i is still the minimal integer satisfying (2.10).

Part 3: If $\mathrm{slink}_*(T) = T$, the result is trivial. Assume otherwise. We start by showing that slink_* does not change j. By (2.11) and (2.10),

$$\beta_j(\mathrm{slink}_*(T)) = \beta_i(T) + j - i > \beta_i(T) \geq \beta_{j-1}(T) \geq \mu_j(T), \tag{2.15}$$

so the j^{th} run of $\mathrm{slink}_*(T)$ must have cells below the j^{th} row. That is, the first j runs of $\mathrm{slink}_*(T)$ do not form a superstandard tableau. To prove that the first $j-1$ runs of $\mathrm{slink}_*(T)$ form a superstandard tableau, consider two cases.

First, suppose that slink_* increases the size of the i^{th} run of T. In this case, we need to show that there are no new cells of the i^{th} run strictly below the i^{th} row, guaranteeing that the first $j-1$ runs comprise a superstandard tableau. Suppose the contrary. By the definition of slink_*, the i^{th} run of $\mathrm{slink}_*(T)$ would have to contain $\mu_i(T)$ cells in the i^{th} row as well as at least one more cell in a lower row. Using this fact and (2.11),

$$\mu_i \leq \beta_i(\mathrm{slink}_*(T)) - 1 = \beta_j(T) + (i-1) - j. \tag{2.16}$$

In particular, i was not chosen minimally, providing the desired contradiction.

Second, suppose slink_* decreases the size of the i^{th} run. The first $j-1$ runs of $\mathrm{slink}_*(T)$ are then the result of removing cells from the i^{th} run of a superstandard tableaux. Applying Part 1 of Lemma 2.4, this must still be a standard Young tableaux. Hence, the first $j-1$ runs of $\mathrm{slink}_*(T)$ form a standard Young tableaux while the first j runs do not. That is, slink_* preserves j.

We still need to show that slink_* does not change i. We proceed by checking that i is the minimal number satisfying (2.10) for $\mathrm{slink}_*(T)$. Since slink_* preserves j and permutes the cells of the first j runs, it must also preserve $\mu(T)$. Further, $\mu_{i+1}(T)$ must be weakly less than β_i, since the first $j-1$ runs of T comprise a superstandard tableau. Hence

$$\mu_{i+1}(\mathrm{slink}_*(T)) = \mu_{i+1}(T) \leq \beta_i(T) = \beta_j(\mathrm{slink}_*(T)) + i - j. \tag{2.17}$$

To show that i is still minimal, we proceed by contradiction. Suppose that $i-1$ also satisfies (2.10) for $\mathrm{slink}_*(T)$. Then

$$\begin{aligned}
\mu_i(T) \geq& \beta_i(T) = \beta_j(\mathrm{slink}_*(T)) + i - j \\
>& \beta_j(\mathrm{slink}_*(T)) + (i-1) - j \geq \mu_i(\mathrm{slink}_*(T)) \\
\geq& \beta_i(\mathrm{slink}_*(T)) = \beta_j(T) + i - j.
\end{aligned} \tag{2.18}$$

Hence, i was not the minimal choice satisfying (2.10) for T. This contradiction concludes the proof of Part 3. $\qquad\square$

Proposition 2.7. *For any $T \in \mathrm{SYT}(n)$, $\mathrm{slink}_* \circ \mathrm{slink}_*(T) = T$.*

10

Proof. If slink$_*$ acts as the identity, then the result it trivial, so assume otherwise. Part 3 of Lemma 2.6 guarantees that slink$_*$ fixes i and j in Definition 2.3. The permutation of runs described in (2.11) is thus an involution, completing the proof. $\square$

Lemma 2.8. *For any* $T \in \mathrm{SYT}(n)$, $\mathrm{slink} \circ \mathrm{slink}_*(T) = \mathrm{slink}_* \circ \mathrm{slink}(T)$.

Proof. If $\mathrm{slink}(T) = T$, then $\mathrm{slink}_*(T) = T$, and the proof is trivial. We may then assume that $\mathrm{slink}(T) \neq T$, and let i and j be as in the definition of $\mathrm{slink}_*(T)$. If $i = j - 1$, then $\mathrm{slink}_*(T) = \mathrm{slink}(T)$ by Part 3 of Lemma 2.4, so $\mathrm{slink}_* \circ \mathrm{slink}(T) = \mathrm{slink}_* \circ \mathrm{slink}_*(T) = T$, by Proposition 2.7. By Part 3 of Lemma 2.6, slink$_*$ preserves the i and j values from Definition 2.3. Hence, $\mathrm{slink} \circ \mathrm{slink}_*(T) = \mathrm{slink}_* \circ \mathrm{slink}_*(T) = T$, completing the $i = j - 1$ case.

Now assume that $i < j - 1$. Applying Lemma 2.6 to $\mathrm{slink} \circ \mathrm{slink}_*(T)$, slink$_*$ effects the i and j^{th} runs, and then slink effects the j and $j - 1^{th}$ runs. Conversely, in $\mathrm{slink}_* \circ \mathrm{slink}(T)$, slink effects the $j - 1$ and j^{th} runs, and then slink$_*$ effects the i and $j - 1^{th}$ runs. In particular, both $\mathrm{slink} \circ \mathrm{slink}_*$ and $\mathrm{slink}_* \circ \mathrm{slink}$ only effects the cells of the $i, j - 1$, and j^{th} runs. By (2.9) and (2.11), the action on $\beta(T)$ gives

$$\begin{aligned} \beta_{i,j-1,j}(\mathrm{slink} \circ \mathrm{slink}_*(T)) &= \beta_{i,j-1,j}(\mathrm{slink}_* \circ \mathrm{slink}(T)) \\ &= (\beta_j(T) + i - j, \beta_i(T) + j - i - 1, \beta_{j-1}(T) + 1). \end{aligned} \tag{2.19}$$

Again, $\mathrm{slink} \circ \mathrm{slink}_*(T)$ and $\mathrm{slink}_* \circ \mathrm{slink}(T)$ only differ in the locations of the i, $j - 1$, and j^{th} runs, but the number of cells in these runs are equal. It thus suffices to show that the i^{th} and j^{th} runs are in the same cells, forcing the $j - 1^{th}$ run to also agree. By Lemma 2.6, the first i runs of both form superstandard tableaux. Hence, the locations of the i^{th} run is fully determined by its inverse descent set, and so must be the same. Now let a be the number of cells in the j^{th} run of T that are in the j^{th} row of T. Applying Definition 2.3, in both $\mathrm{slink} \circ \mathrm{slink}_*(T)$ and $\mathrm{slink}_* \circ \mathrm{slink}(T)$, the j^{th} run will consist of the a cells in the j^{th} row and the last $\beta_{j-1}(T) + 1 - a$ cells of the j^{th} run of T. Hence, both tableaux have the j^{th} runs in the same cells, completing the proof. $\square$

Proposition 2.9. *Let* $T \in \mathrm{SYT}(n)$. *Then* $\mathrm{slink}^{j-i}(T) = \mathrm{slink}^{j-i-1} \circ \mathrm{slink}_*(T)$, *where i and j are as in Definition 2.3. In particular,* $\equiv^0$ *is a refinement of* $\equiv^1$.

Proof. We will prove the first part by induction on $j - i$. If j does not exist, then the result is trivial. If $j - i = 1$, then $\mathrm{slink}(T) = \mathrm{slink}_*(T)$, by Part 3 of Lemma 2.4. Now assume the result whenever $j - i \leq k$, and assume $j - i = k + 1 > 1$.

We would like to apply our inductive hypothesis to $\mathrm{slink}(T)$. To justify this, recall that Part 2 of Lemma 2.6 guarantees that slink fixes the i value of T while lowering the j value by 1, hence lowering $j - i$ by 1. Applying our inductive hypothesis, as well as Lemma 2.8, gives

$$\mathrm{slink}^{j-i}(T) = \mathrm{slink}^{j-i-1} \circ \mathrm{slink}(T) = \mathrm{slink}^{j-i-2} \circ \mathrm{slink}_* \circ \mathrm{slink}(T) = \mathrm{slink}^{j-i-1} \circ \mathrm{slink}_*(T), \tag{2.20}$$

completing the proof of the first part of the lemma. The second part of the proposition now follows from Definition 2.5, since $\mathrm{slink} *(T)$ is always connected to T by some sequence of applications of slink. $\square$

Lemma 2.10. *Given any partition λ and tableau $T \in \mathrm{SYT}(\lambda)$, exactly one of the following holds.*

1. $T = U_\lambda$.

2. $s_{\beta(T)} = -s_{\beta(\mathrm{slink}_*(T))}$

Proof. First, we show that the two cases of the lemma are disjoint. If T is described by Case 1, then $\mathrm{slink}_*(T) = T$ and $s_{\beta(T)} = s_\lambda \neq 0$. Hence, $s_{\beta(T)} \neq -s_{\beta(\mathrm{slink}_*(T))}$, and Case 2 does not apply.

It now suffices to assume that T is not described by Case 1 and show that it satisfies Case 2. By Proposition 2.9, T is connected to $\mathrm{slink}_*(T)$ by an odd number of applications of slink. Applying Part 2 of Lemma 2.4, $s_{\beta(T)} = -s_{\beta(\mathrm{slink}_*(T))}$. $\square$

Lemma 2.11. *If $\mathcal{C}$ is any equivalence class of $\equiv^0$ or $\equiv^1$ such that $\mathcal{C} \neq \{U_\lambda\}$ for any partition λ, then $\sum_{T \in \mathcal{C}} s_{\beta(T)} = 0$.*

Proof. We proceed by providing a sign reversing involution. First notice that U_λ is always alone in its $\equiv^0$ and $\equiv^1$ classes, so we may assume $U_\lambda \notin \mathcal{C}$. Each $T \in \mathcal{C}$ contributes one composition Schur function to the sum, which is in turn equal to 0 or $\pm s_\lambda$, for some $\lambda \vdash n$. By Proposition 2.9, slink_* is an involution on $\mathcal{C}$. By Lemma 2.10 and the assumption that $\mathcal{C} \neq \{U_\lambda\}$, $s_{\beta(T)} + s_{\beta(\mathrm{slink}_*(T))} = 0$. Thus, the sum of $s_{\beta(T)}$ over $T \in \mathcal{C}$ must equal zero. $\square$

Next, we add another equivalence relation, which we term *restricted dual equivalence.*

Definition 2.12. For any $T \in \mathrm{SYT}(n)$ and any $i \in [2, n-2]$, let

$$d_i^R(T) = \begin{cases} T & i+1 \in \mathrm{ID}(T) \cap \mathrm{ID}(d_i(T)), \\ d_i(T) & \text{otherwise.} \end{cases}$$

Further, let $\equiv^2$ be the equivalence relation on standard Young tableaux generated by the action of all d_i^R.

Theorem 2.13. *The equivalence relation $\equiv^0$ is a refinement of $\equiv^1$, which in turn refines $\equiv^2$.*

Proof. The first part of the statement follows directly from Proposition 2.9, so we need only show that $\equiv^1$ refines $\equiv^2$. It suffices to show that the action of slink on $T \in \mathrm{SYT}(n)$ can always be achieved by a series of d_i^R.

First consider the case where $T \in \mathrm{SYT}(\lambda)$ and λ has at most two rows. It follows directly from the Definition 2.3 that slink fixes the cell containing n, which is the only property we will need. In the two row case, $\{i, i+1\}$ cannot be a subset of $\mathrm{ID}(T)$, because $i+2$ would have to be two rows above i. Direct inspection of the definition of d_i or (2.3) shows that $i+1 \in \mathrm{ID}(T) \cap \mathrm{ID}(d_i(T))$ implies $\{i, i+1\} \subset \mathrm{ID}(T)$ or $\{i, i+1\} \subset \mathrm{ID}(d_i(T))$. Thus, $d_i^R = d_i$ whenever $i \neq n-1$. By Lemma 2.1, d_i acts transitively on standard Young tableaux of a fixed shape, so d_i^R acts transitively on the set of tableaux in $\mathrm{SYT}(\lambda)$ with n in a fixed cell. Thus, some series of d_i^R yields the action of slink on T.

12

Next, suppose that T has more than two rows. If $\mathrm{slink}(T) = T$, the result is trivial. Otherwise, slink acts by changing the $j - 1$ and j^{th} runs. Let $[a, b]$ be the values in these runs, and consider the action of d_i^R on $\mathrm{fl}(\mathrm{rw}(T)|_{[a,b]})$. In particular, this is the row reading word of some two row standard Young tableau. Applying the two row case, we are free to apply slink to this two row tableau via some series of d_i^R. Adding $a - 1$ to each i in this sequence provides the sequence of d_i^R to apply to T in order to achieve $\mathrm{slink}(T)$. □

Lemma 2.14. *If $\mathcal{C}$ is any union of equivalence classes of $\equiv^k$ with $k = 0, 1, 2$, then*

$$\sum_{T \in \mathcal{C}} s_{\beta(T)} = \sum_{U_\lambda \in \mathcal{C}} s_\lambda.$$

Proof. Applying Theorem 2.13, it suffices to only consider $\equiv^0$. By Lemma 2.11, we may further omit any equivalence classes not of the form $\{U_\lambda\}$. Thus,

$$\sum_{T \in \mathcal{C}} s_{\beta(T)} = \sum_{U_\lambda \in \mathcal{C}} s_{\beta(U_\lambda)} = \sum_{U_\lambda \in \mathcal{C}} s_\lambda. \tag{2.21}$$

□

Proof of Theorem 1.1. For $k = 0, 1, 2$, let $\mathcal{C}$ be the disjoint union of equivalence classes of $\equiv^k$ such that $f = \sum_{T \in \mathcal{C}} F_{ID(T)}$ is a symmetric function. By Lemmas 2.2 and 2.14,

$$f = \sum_{T \in \mathcal{C}} s_{\beta(T)} = \sum_{\lambda \vdash n} c_\lambda s_\lambda. \tag{2.22}$$

where c_λ is the multiplicity of U_λ in $\mathcal{C}$. □

2.3 The equivalence relations on permutations

We extend these results to permutations via the RSK correspondence. Recall that a function defined on standard Young tableaux acts on permutations via insertion tableaux by using the RSK correspondence, acting on the insertion tableaux and fixing the recording tableaux.

Definition 2.15. Let $\mathrm{slink}, \mathrm{slink}_*$, and d_i^R act on permutations via insertion tableaux. For $k = 0, 1, 2$, let $\pi \equiv^k \pi'$ if $P(\pi) \equiv^k P(\pi')$ and $Q(\pi) = Q(\pi')$.

Notice that, $\equiv^k$ is defined so that the same function generates equivalence classes on permutations as on tableaux. In light of (2.4) and (2.5), $\equiv^2$ is defined on permutations by the natural action of d_i^R on permutations. That is, $d_i^R(\pi) = d_i(\pi)$ unless $i + 1 \in ID(\pi) \cap ID(d_i(\pi))$, in which case d_i^R acts trivially. By considering the action of d_i^R on $\mathrm{fl}(\pi|_{[i-1,i+2]})$, we may express all nontrivial actions as

$$x1y4 \quad x14y \quad x41y \quad x3y4 \quad x34y, \tag{2.23}$$

where d_i^R swaps the values x and y in $\mathrm{fl}(\pi|_{[i-1,i+2]})$ while fixing values of π not in $[i - 1, i + 1]$.

Proposition 2.16. *The action of K_j on permutations commutes with d_i^R, slink, and slink_*.*

Proof. To allow us to consider all three cases at once, let $\phi = d_i^R$, slink, or slink$_*$. For any $\pi \in S_n$, it suffices to consider the effect of each function in question on $(P(\pi), Q(\pi))$. Recall that K_j acts only on $Q(\pi)$, and this action is independent of $P(\pi)$. Similarly, ϕ acts only $P(\pi)$, and this action is independent of $Q(\pi)$. Hence $K_j \circ \phi$ and $\phi \circ K_j$, and both act via the RSK correspondence to yield the unique permutation sent to $(P(\phi(\pi)), Q(K_j(\pi)))$, completing the proof. □

Remark 2.17. The action of Knuth equivalence is often introduced in relation to the process of jeu de taquin on skew tableaux. Had we taken this approach, the translated result would state that the action of slink, slink$_*$, and d_i^R on standard skew tableaux commutes with jeu de taquin.

Corollary 2.18. *For any $k = 0, 1, 2$, let $\mathcal{C} \subset S_n$ be a union of equivalence classes of $\equiv^k$ such that $f = \sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\pi)}$ is a symmetric function. Then*

$$f = \sum_{\lambda \vdash n} c_\lambda s_\lambda,$$

where $c_\lambda = |\{\pi \in \mathcal{C} \colon \pi \in \mathrm{SYam}(\lambda)\}|$.

Proof. Each $\equiv^k$ equivalence class on permutations is sent to an $\equiv^k$ equivalence class on tableaux by the function P, while preserving inverse descent sets. Consider the disjoint union of these classes of tableaux. By Theorem 1.1, $f = \sum c_\lambda s_\lambda$, where c_λ is the multiplicity of U_λ in this disjoint union. The set $\mathrm{SYam}(\lambda)$ is defined to be the set of permutations sent to U_λ by P. Thus, $c_\lambda = |\{\pi \in \mathcal{C} \colon \pi \in \mathrm{SYam}(\lambda)\}|$. □

We end this section by considering generating functions of $\equiv^k$ classes.

Definition 2.19. *For $k = 0, 1, 2$, let $\{f^{(k)}\}$ be the set of functions that can be realized as $\sum_{T \in \mathcal{C}} F_{\mathrm{ID}(T)}$, where $\mathcal{C}$ is a single equivalence class of $\equiv^k$.*

We use the notation $\{f^{(k)}\}$ because it is currently unknown how to best index these sets. In the next section, we will show that each $\{f^{(k)}\}$ forms a spanning set of the quasisymmetric functions.

Conjecture 2.20. *The set $\{f^{(2)}\}$ forms a basis for the quasisymmetric functions.*

This conjecture has been verified up to degree 11 with the aid of a computer.

3   Extending to the quasisymmetric Schur functions

We turn our attention to the combinatorics of quasisymmetric Schur functions, placing them near the top of our poset of quasisymmetric functions.

3.1 preliminaries

A *standard reverse composition tableau* T is a standard filling of a composition shape that is increasing down the first column, decreasing across rows from left to right, and that satisfies the following additional property. Given any two values $a > b$ that are adjacent in a row of T, no value in (b, a) may occur below and in the same column as b. If a is the furthest right value in its row, treat the entry immediately to its right as 0. The set of standard reverse composition tableaux is denoted SRCT(α). The *bent reading word* of $T \in$ SRCT(α), bw(T), is given by reading down columns from right to left and then up the leftmost column. Examples are given in Figure 7. We may now the *quasisymmetric Schur functions* by

$$\mathcal{S}_\alpha = \sum_{T \in \mathrm{SRCT}(\alpha)} F_{\mathrm{ID}(\mathrm{bw}(T))}. \tag{3.1}$$

Importantly, $\{\mathcal{S}_\alpha\}$ is a basis for the quasisymmetric functions. For a more thorough treatment of the quasisymmetric Schur functions, see [Luoto et al., 2013].

Remark 3.1. Traditionally, the inverse descent set of $T \in$ SRCT(α) is the set of i such that $i + 1$ is in a column weakly to the right of i. This is equivalent to defining the inverse descent set via the bent reading word.

A *standard reverse Young tableau* of partition shape λ is a standard filling of shape λ that is decreasing up columns and across rows from left to right. The set of all standard reverse Young tableau is denoted SRT(λ). The *reverse column word* of $T \in$ SRT(λ), cw(T), will be the reverse of the normal convention, proceeding up columns, right to left. It can then be shown (see [Luoto et al., 2013, Ch. 3.2]) that the basis of Schur functions may be rewritten as

$$s_\lambda = \sum_{T \in \mathrm{SRT}(\lambda)} F_{\mathrm{ID}(\mathrm{cw}(T))}. \tag{3.2}$$

The set $\mathcal{C}(\alpha) \subset \mathrm{SRT}(\lambda(\alpha))$ is given by applying Mason's bijection to SRCT(α), sorting the columns, and bottom justifying. We denote this map by ρ. Importantly, ρ preserves the set of values in each column. Further, it preserves inverse descent sets, since both SRT and SRCT have the property that i is an inverse descent of a filling if and only if $i + 1$ is in a column weakly to its right. Thus,

$$\mathcal{S}_\alpha = \sum_{T \in \mathcal{C}(\alpha)} F_{\mathrm{ID}(T)}. \tag{3.3}$$

The map ρ also has an inverse defined as follows. Working one column at a time from left to right, start with the second column. Now consider each value x in the column in order from greatest to least. Out of all remaining locations in the column, place x to the right of highest cell whose value is greater than x. Figure 7 gives examples, while the original definition is in [Mason, 2008].
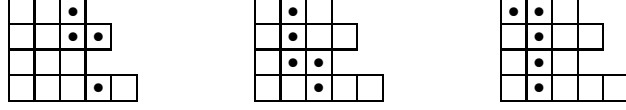
Figure 6: Three α-pistols of $\alpha = (5, 3, 4, 3)$.

3.2   Transitive actions on $\mathrm{SRCT}(\alpha)$

For any cell c of α, the set of cells weakly below c in its column or weakly above c in the column to its left is called an α-*pistol*, as in Figure 3.2.

Let $\tilde{d}_i : S_n \to S_n$ be the involution that cyclically permutes the values $i - 1, i,$ and $i + 1$ as

$$\tilde{d}_i(\ldots i \ldots i - 1 \ldots i + 1 \ldots) = (\ldots i - 1 \ldots i + 1 \ldots i \ldots),$$
$$\tilde{d}_i(\ldots i \ldots i + 1 \ldots i - 1 \ldots) = (\ldots i + 1 \ldots i - 1 \ldots i \ldots), \tag{3.4}$$

and that acts as the identity if i occurs between $i - 1$ and $i + 1$. For example, $\tilde{d}_3 \circ \tilde{d}_2(4123) = \tilde{d}_3(4123) = 3142$. Let d_i and $\tilde{d}_i$ act on $\mathrm{SRCT}(\alpha)$ via the bent reading word.

We use d_i and $\tilde{d}_i$ to define an involution on $\mathrm{SRCT}(\alpha)$.

$$D_i^Q(T) = \begin{cases} T & \text{if } i - 1, \ i, \text{ and } i + 1 \text{ are contained in two cells of} \\ & \text{the first column and one in the second column,} \\ \tilde{d}_i(T) & \text{else, if } i - 1, \ i, \text{ and } i + 1 \text{ are contained in an } \alpha\text{-pistol,} \\ d_i(T) & \text{otherwise.} \end{cases} \tag{3.5}$$

As an example, we may take $\mathrm{bw}(T) = \pi = 3265874$ and $\alpha = (3, 2, 3)$ as in Figure 7. Examples of each of the three cases are as follows. First, $D_7^Q(T) = T$. Then, $D_6^Q(T) = \tilde{d}_6(T)$, which has bent reading word 3257468. Finally, $D_3^Q(T) = d_3(T)$, which has bent reading word 4265873.
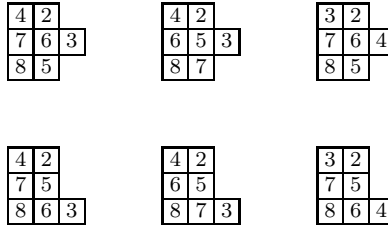

Figure 7: Three standard reverse composition tableaux of shape $\alpha = (2, 3, 2)$ above their image via Mason's bijection. From the left, $T \in \mathrm{SRCT}(\alpha)$ with bent reading word $\pi = 3265874$, followed by $D_6^Q(T)$, and then $D_3^Q(T)$.

Remark 3.2. The function D_i^Q is closely related to its namesake, D_i, defined in [Assaf, 2015].

Proposition 3.3. *The set of involutions D_i^Q preserves and acts transitively on $\mathrm{SRCT}(\alpha)$.*

Proof. The first part, that D_i^Q sends SRCT(α) to itself, is a straight forward matter of considering the possible arrangements of $[i-1, i+1]$, and it is recommended for the reader. We will instead focus on proving the transitivity of the action of D_i^Q on SRCT(α).

The case where $\alpha \vDash 1$ is trivial, so we may proceed by induction, letting $\alpha \vDash n$ and assuming the result for compositions of $n-1$. It suffices to show that the value 1 may be moved from the top row to any other possible location of 1. From there, the inductive hypothesis completes the claim. We will consider four cases, as displayed in Figure 8.

First suppose that we wish to move 1 to a column strictly to the left of its position in the top row. Place the 2 in the desired cell and place 3 directly to the left of 1. In this case, D_i^Q acts via d_i, swapping the 1 with the 2.

Second, suppose we wish to move 1 to a cell below it and in the same column. Place the 3 in the desired cell and place 2 directly to the left of 1. In this case D_i^Q acts via $\tilde{d}_i$, cycling the three numbers.

Next, notice that no cell a single column to the right of 1 is an allowable cell to place 1. This follows by considering the triple formed by 1's original location, the lack of a cell immediately to the right of that, and any desired new location.

Finally, suppose we wish to move 1 at least two columns to the right. Applying the second case, we may move 1 as far down as possible within its own column. Then place the 2 in the desired cell, and place 3 in any allowable location between the 1 and the 2 in the bent reading word. In this fourth case, D_i^Q acts via d_i, swapping the 1 with the 2. $\square$

Figure 8: The four cases, in order, demonstrating the transitivity of the action of D_i^Q on SRCT(α).

We want to express the action of D_i^Q on $\mathcal{C}(\alpha)$ via Mason's bijection. Given $T \in$ SRT(λ), define the *elementary quasi-dual equivalence relations* as

$$d_i^Q(T) = \begin{cases} T & \text{if } i-1, i, \text{ and } i+1 \text{ are in two cells of the first} \\ & \text{column and one in the second column,} \\ d_i(T) & \text{otherwise.} \end{cases} \tag{3.6}$$

As with earlier involutions, we define the action of d_i^Q on a permutation via insertion tableaux.

Proposition 3.4. *For any $T \in$ SRCT(α),*

$$d_i^Q \circ \rho(T) = \rho \circ D_i^Q(T)$$

Proof. First notice that each $S \in$ SRT(λ) and each $T \in$ SRCT(α) are completely determined by the list of values that belong in each column. In the latter case, this is made clear by the algorithm for ρ^{-1}. Because ρ preserves the list of values in each column, it suffices to show that D_i^Q and d_i^Q have the same effect on the values in the columns of T and $\rho(T)$, respectively.

Consider the case where D_i^Q acts trivially on T. Since ρ preserves inverse descent sets, d_i and $\tilde{d}_i$ act trivially on T if and only if they act trivially on $\rho(T)$. Furthermore, ρ preserves the sets of values in each column, so D_i^Q acts trivially on T if and only if d_i^Q acts trivially on $\rho(T)$.

Assume that D_i^Q and d_i^Q do not act trivially. In particular, $i-1, i$, and $i+1$ may not all be in the same column, since this would require i to be between $i-1$ and $i+1$ in the reading word. If $i-1, i$, and $i+1$ are all in different columns, then d_i^Q and D_i^Q will both act by swapping the two outside values, leaving the value in the middle column in place. Thus, it suffices to consider the case where $i-1, i$, and $i+1$ are in exactly two columns.

Without loss of generality, assume that i shares its column with either $i-1$ or $i+1$ — if this were not the case, we could consider $D_i^Q(T)$ instead. Since d_i^Q and D_i^Q always change the column of i, they must both move i into the same column, forcing $i-1$ and $i+1$ to share the remaining column. Since D_i^Q and d_i^Q have the same action on the values in each column of T and $\rho(T)$, respectively, we must have $d_i^Q \circ \rho(T) = \rho \circ D_i^Q(T)$. $\square$

Corollary 3.5. *The set of d_i^Q preserve and act transitively on $\mathcal{C}(\alpha)$.*

Proof. The statement follows from Lemma 3.3 and Poposition 3.4. $\square$

Remark 3.6. In [Bessenrodt et al., 2011], there is a similar analysis of the relationship between ρ and dual equivalence on SRT for the sake of proving a quasisymmetric Littlewood-Richardson rule. We have made these relationships more explicit by introducing the functions d_i^Q and D_i^Q.

Recall from Definition 2.19 that $\{f^{(k)}\}$ is the set of functions of the form $f = \sum_{T \in C} F_{\mathrm{ID}(T)}$, where $\mathcal{C}$ is any $\equiv^k$ class. Each $\equiv^k$ also defines relations on $\mathrm{SRT}(\lambda)$ via the reverse column reading word. In this way, it makes sense to compare the equivalence classes generated by d_i^Q to the classes of $\equiv^k$ and to compare $\{\mathcal{S}_\alpha\}$ to $\{f^{(k)}\}$.

Theorem 3.7. *Quasi-dual equivalence is a coarsening of $\equiv^k$, for each $k = 0, 1, 2$. In particular, each quasisymmetric Schur function is a nonnegative sum of functions in $\{f^{(k)}\}$, for each $k = 0, 1, 2$.*

Proof. In light of Theorem 2.13, it suffices to consider $\equiv^2$. Notice that d_i^R and d_i^Q have the same definition when they act nontrivially on some $T \in \mathrm{SRT}(\lambda)$. It then suffices to show that when d_i^Q acts trivially on $T \in \mathrm{SRT}(\lambda)$, so does d_i^R. If d_i acts trivially, then the result is immediate, so assume otherwise. We may thus assume that two values of $[i-1, i+1]$ are in the first column of T. Notice that the action of d_i will not change the number of values in $[i-1, i+1]$ that are in the first column and that $i \notin \mathrm{ID}(T)$ implies that i is not in the first column. Hence, if two values of $[i-1, i+1]$ are in the first column of T, then $|\mathrm{ID}(T) \cap [i-1, i+1]| = |\mathrm{ID}(d_i(T)) \cap [i-1, i+1]| \geq 2$. Comparing with (2.23) shows that this never happens when d_i^R acts nontrivially. Hence, d_i^R must act trivially. This completes the proof. $\square$

Remark 3.8. The previous proof only used the fact that d_i^R acts trivially when two values in $[i-1, i+1]$ are in the first column, which is less strict than the requirement on d_i^Q. If we wished, we could add yet another equivalence relation to our poset by using this intermediate condition.

Corollary 3.9. *The involution d_i^R preserves $\mathcal{C}(\alpha)$.*

Proof. The corollary is a consequence of Theorem 3.7. □

As a side note, the set of d_i^R do not always act transitively on $\mathcal{C}(\alpha)$, as can be seen by inspecting $\alpha = (2, 2, 2)$.

Corollary 3.10. *For each $k = 0, 1, 2$, the set of functions $\{f^{(k)}\}$ of the form $f = \sum_{T \in C} F_{\mathrm{ID}(T)}$ forms a spanning set of the quasisymmetric functions.*

Proof. By Theorem 3.7, all quasisymmetric Schur functions must be positive sums of functions in $\{f^{(k)}\}$. The quasisymmetric Schur functions form a basis for the quasisymmetric functions, and so each $\{f^{(k)}\}$ must form a spanning set. □

We also have a new proof of the following, which is stated in [Bessenrodt et al., 2014, Lemma 2.21] as a corollary to the decomposition of Schur functions into quasisymmetric Schur functions in [Haglund et al., 2011].

Corollary 3.11. *Any symmetric function that is a positive sum of quasisymmetric Schur functions is Schur positive.*

Proof. This follows from Theorem 1.1 and Theorem 3.7. □

4 Extending to shifted dual equivalence

4.1 preliminaries

A partition λ is termed a *strict partition* if $\lambda = (\lambda_1 > \lambda_2 > \ldots > \lambda_k > 0)$, the *shifted shape* λ is the set of squares in positions $\{(i, j)$ given $1 \leq i \leq k, i \leq j \leq \lambda_i + i - 1\}$. A *standard shifted tableau* T is a standard filling of a shifted shape that is increasing up columns and across rows from left to right. For example, see Figure 9. Let $\mathrm{SST}(\lambda)$ be the set of standard shifted tableaux of shifted shape λ.

$$T = \begin{array}{c}\boxed{7}\\ \boxed{5}\,\boxed{6}\\ \boxed{1}\,\boxed{2}\,\boxed{3}\,\boxed{4}\,\boxed{8}\end{array}, \qquad h_3(T) = \begin{array}{c}\boxed{7}\\ \boxed{4}\,\boxed{6}\\ \boxed{1}\,\boxed{2}\,\boxed{3}\,\boxed{5}\,\boxed{8}\end{array}, \qquad h_5(T) = \begin{array}{c}\boxed{8}\\ \boxed{5}\,\boxed{6}\\ \boxed{1}\,\boxed{2}\,\boxed{3}\,\boxed{4}\,\boxed{7}\end{array}.$$

Figure 9: Three shifted dual equivalent standard shifted tableau of shape $\lambda = (5, 2, 1)$.

For a treatment on shifted Knuth equivalence, shifted dual equivalence, and shifted jeu de taquin, see [Sagan, 1987] and [Haiman, 1992]. For our purposes, it suffices to define a shifted analog of dual equivalence.

Definition 4.1 ([Haiman, 1992]). *Given a permutation $\pi \in S_n$, define the* elementary shifted dual equivalence h_i *for all $1 \leq i \leq n - 3$ as follows. If $n \leq 3$, then $h_1(\pi) = \pi$. If $n = 4$, then $h_1(\pi)$ acts by swapping x and y in the cases below,*

$$1x2y \quad x12y \quad 1x4y \quad x14y \quad 4x1y \quad x41y \quad 4x3y \quad x43y, \tag{4.1}$$

19

and $h_1(\pi) = \pi$ otherwise. If $n > 4$, then h_i is the involution that fixes values not in $I = [i, i+3]$ and permutes the values in I via $\mathrm{fl}(h_i(\pi)|_I) = h_1(\mathrm{fl}(\pi|_I))$.

As an example, $h_1(24531) = 14532$, $h_2(25134) = 24135$, and $h_3(314526) = 314526$. Notice that every nontrivial action of h_i is equal to either d_{i+1} or d_{i+2}. Two permutations are *shifted dual equivalent* if they are connected by some sequence of elementary shifted dual equivalences.

Given a standard shifted tableau T, we define $h_i(T)$ as the result of letting h_i act on the row reading word of T. Observe that $h_i(T)$ is also a standard shifted tableau. We can define an equivalence relation on standard shifted tableaux by saying T and $h_i(T)$ are *shifted dual equivalent* for all i. See Figure 9 for an example.

Theorem 4.2 (Prop. 2.4, [Haiman, 1992]). *For all strict partitions λ, the set of h_i act transitively on $\mathrm{SSYT}(\lambda)$.*

As with $\{f^{(k)}\}$, we wish to consider the generating functions of shifted dual equivalence.

Definition 4.3. *Let $\{f^{(h)}\}$ be the set of functions that can be realized as $\sum_{T \in \mathcal{C}} F_{\mathrm{ID}(T)}$, where $\mathcal{C}$ is a single shifted dual equivalence class of permutations.*

To study $\{f^{(h)}\}$, we will need some intermediary involutions. For an inverse descent set $D \subset [n-1]$, let D^C be the compliment of D in $[n-1]$. The involution on quasisymmetric functions that sends F_D to F_{D^C} and s_λ to s'_λ is denoted ω. Given $\pi \in S_n$, let the *reverse* of π, $\mathrm{rev}(\pi)$, be the result of reading π from right to left. Let the *flip* of π, $\mathrm{flip}(\pi)$, be the result of taking $\mathrm{rev}(\pi)$ and inverting all values, sending i to $n + 1 - i$. For example, $\mathrm{rev}(21435) = 53412$ and $\mathrm{flip}(21435) = 13254$. Notice that $\mathrm{ID}(\mathrm{rev}(\pi)) = \mathrm{ID}(\pi)^C$, and $\mathrm{ID}(\mathrm{flip}(\pi))$ is the result of sending each $i \in \mathrm{ID}(\pi)$ to $n - i$. Thus, if $f = \sum_\pi F_{\mathrm{ID}(\pi)}$, then

$$\sum_\pi F_{\mathrm{ID}(\mathrm{rev}(\pi))} = \sum_\pi F_{\mathrm{ID}(\pi)^C} = \omega(f). \tag{4.2}$$

Further, if $s_\lambda = \sum_\pi F_{\mathrm{ID}(\pi)}$, then

$$s_{\lambda'} = \sum_\pi F_{\mathrm{ID}(\mathrm{rev}(\pi))} \quad \text{and} \quad s_\lambda = \sum_\pi F_{\mathrm{ID}(\mathrm{flip}(\pi))}. \tag{4.3}$$

Though (4.3) is well known, we can sketch a proof as follows. By considering the equation for s_λ in (1.1), we may send each $T \in \mathrm{SYT}(\lambda)$ to its conjugate, sending the values in the i^{th} row λ to the i^{th} column of λ'. The resulting filling in $\mathrm{SYT}(\lambda')$ will have the same inverse descent set as $\mathrm{rev}(\mathrm{rw}(T))$. If we instead invert every value in T, we may appeal to (3.2) by sending each value i to $n - i$. We can then use the reverse column reading word to give a filling in $\mathrm{SRT}(\lambda)$ with the same inverse descent set as $\mathrm{flip}(\mathrm{rw}(T))$.

4.2 From restricted dual equivalence to shifted dual equivalence

Proposition 4.4. *Given distinct $\pi, \pi' \in S_n$, if $d_i^R(\mathrm{rev}(\pi)) = \mathrm{rev}(\pi')$, then $h_{i-1}(\pi) = \pi'$. If $d_i^R(\mathrm{flip}(\pi)) = \mathrm{flip}(\pi')$, then $h_{n-i}(\pi) = \pi'$.*

20

Proof. Taking the reverse of each nontrivial action of d_i^R, as expressed in (2.23), gives

$$x14y \quad x41y \quad 4x1y \quad 4x3y \quad x43y, \tag{4.4}$$

each of which is present in (4.1). Since h_i acts on $[i, i+3]$ and d_i^R acts on $[i-1, i+2]$, the action of d_i^R is achieved by h_{i-1}, completing the first part of the proposition.

Taking the flip of each nontrivial action of d_i^R, as expressed in (2.23), gives

$$x41y \quad x14y \quad 1x4y \quad 1x2y \quad x12y, \tag{4.5}$$

each of which is present in (4.1). Taking the flip of the word sends i to $n+1-i$. Again, h_i acts on $[i, i+3]$ while d_i^R acts via the flip of π on $[n-i-1, n-i+2]$, so d_i^R is achieved by h_{n-i-1}. This completes the proof of the second part of the proposition. $\square$

Theorem 4.5. *For $k = 0, 1, 2$, the relation $\equiv^k$ on S_n is a refinement of shifted dual equivalence on the reverse of the permutations in S_n. In particular, every function in $\{f^{(h)}\}$ is a nonnegative sum of functions in $\{\omega(f) : f \in \{f^k\}\}$, for each $k = 0, 1, 2$. Further, the relation $\equiv^k$ on S_n is a refinement of the relation generated by shifted dual equivalence on the flip of the permutations in S_n.*

Proof. By Theorem 2.13, it suffices to only consider the $k = 2$ case. The statements about refining equivalence relations are then direct applications of Proposition 4.4 to the definitions of $\equiv^2$ and shifted dual equivalence. The statement about $\{f^h\}$ follows directly from Definitions 2.19 and 4.3, along with (4.2). $\square$

Corollary 4.6. *Let $\mathcal{C}$ be any union of shifted dual equivalence classes of S_n such that $f = \sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\pi)}$ is symmetric. Then*

$$f = \sum_{\lambda \vdash n} |\{\pi \in \mathcal{C} : \mathrm{rev}(\pi) \in \mathrm{SYam}(\lambda)\}| \cdot s_{\lambda'},$$

$$= \sum_{\lambda \vdash n} |\{\pi \in \mathcal{C} : \mathrm{flip}(\pi) \in \mathrm{SYam}(\lambda)\}| \cdot s_{\lambda}.$$

In particular, f is Schur positive.

Proof. By Proposition 4.4, reversing all of the permutations in $\mathcal{C}$ yields some $\mathcal{D}$ that is the union of $\equiv^2$ classes. By Corollary 2.18

$$\sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\mathrm{rev}(\pi))} = \sum_{\lambda \vdash n} |\{\pi \in \mathcal{D} : \pi \in \mathrm{SYam}(\lambda)\}| \cdot s_{\lambda}. \tag{4.6}$$

Reversing each permutation has the effect of conjugating the resulting Schur functions, so

$$\sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\pi)} = \sum_{\lambda \vdash n} |\{\pi \in \mathcal{C} : \mathrm{rev}(\pi) \in \mathrm{SYam}(\lambda)\}| \cdot s_{\lambda'}. \tag{4.7}$$

Similarly, By Proposition 4.4, flipping all of the permutations in $\mathcal{C}$ yields some $\mathcal{D}$ that is the union of $\equiv^2$ classes. By Corollary 2.18

$$\sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\mathrm{flip}(\pi))} = \sum_{\lambda \vdash n} |\{\pi \in \mathcal{D} \colon \pi \in \mathrm{SYam}(\lambda)\}| \cdot s_\lambda. \tag{4.8}$$

Flipping each permutation has no effect on the resulting the Schur functions, so

$$\sum_{\pi \in \mathcal{C}} F_{\mathrm{ID}(\pi)} = \sum_{\lambda \vdash n} |\{\pi \in \mathcal{C} \colon \mathrm{flip}(\pi) \in \mathrm{SYam}(\lambda)\}| \cdot s_\lambda. \tag{4.9}$$

$\square$

Remark 4.7. Traditionally, the generating function over standard shifted tableaux uses peak sets rather than inverse descent sets. Accordingly, it is written as a sum of peak quasisymmetric functions rather than fundamental quasisymmetric functions. It is thus surprising to find a simple condition for Schur positivity with the less studied generating functions of shifted dual equivalence classes given above.

We end by considering when $\equiv^2$ gives an entire shifted dual equivalence class. Put a different way, our final result gives a converse to Proposition 4.4, so long as we only consider the row reading words of shifted tableaux in $\mathrm{SSYT}(\lambda)$.

Proposition 4.8. *For all strict partitions λ, the set of d_i^R acts transitively on $\mathrm{SSYT}(\lambda)$ via the flip of the row reading word.*

Proof. Since d_i^R acts via the flip of a row reading word, it acts via the actions in (4.5). By Theorem 4.2 and Proposition 4.4, this action is an involution on $\mathrm{SSYT}(\lambda)$. To prove transitivity, we proceed by induction on n.

If $n = 1$, then the result is trivial. Assume $n \geq 2$ and that the result holds for strict partitions of size less than n. Consider any $T \in \mathrm{SSYT}(\lambda)$ such that n is in the northeast corner, c_1, in the highest row. It suffices to show that we may move n into each of the lower northeast corners by applying the involutions in (4.5), since we may use the inductive hypothesis to rearrange values in $[n-1]$ as needed. If there are no other northeast corners, we are done. Otherwise, the inductive hypothesis guarantees that $n-1$ may be placed in the desired corner, c_2. Then the values $n-3$ and $n-2$ may be placed in the last of the allowable cells before c_2, in row reading order, as in Figure 10. One of the last two involution in (4.5) then swaps the n and $n-1$, as desired. $\square$

References

S. H. Assaf. Dual equivalence graphs I: A new paradigm for Schur positivity. *arXiv:1506.03798*, 2015.

C. Bessenrodt, K. Luoto, and S. van Willigenburg. Skew quasisymmetric Schur functions and noncommutative Schur functions. *Advances in Mathematics*, 226(5):4492–4532, 2011.
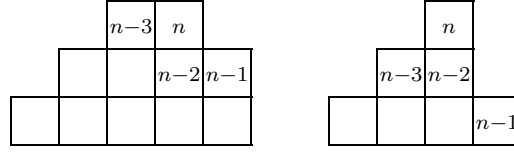
Figure 10: Example placements of the values in $[n-3, n]$ to move n from the top row to a lower row.

C. Bessenrodt, V. V. Tewari, and S. J. van Willigenburg. Symmetric skew quasisymmetric Schur functions. *arXiv preprint arXiv:1410.2934*, 2014.

S. Billey, Z. Hamaker, A. Roberts, and B. Young. Coxeter-Knuth graphs and a signed Little map for type B reduced words. *Electron. J. of Combin.*, 21(P4.6), 2014.

E. Egge, N. A. Loehr, and G. S. Warrington. From quasisymmetric expansions to Schur expansions via a modified inverse Kostka matrix. *European Journal of Combinatorics*, 31(8): 2014–2027, 2010.

I. M. Gessel. Multipartite P-partitions and inner products of skew Schur functions. *Contemp. Math*, 34:289–317, 1984.

J. Haglund, M. Haiman, and N. Loehr. A combinatorial formula for Macdonald polynomials. *J. Amer. Math. Soc*, 18(3), 2005.

J. Haglund, K. Luoto, S. Mason, and S. van Willigenburg. Quasisymmetric Schur functions. *Journal of Combinatorial Theory, Series A*, 118(2):463–490, 2011.

M. D. Haiman. Dual equivalence with applications, including a conjecture of Proctor. *Discrete Math*, 99(1):79–113, 1992.

N. A. Loehr and G. S. Warrington. Quasisymmetric expansion of Schur function plethysms. *Proc. Amer. Math. Soc*, 140(4):1159–1171, 2012.

K. Luoto, S. Mykytiuk, and S. van Willigenburg. *An introduction to quasisymmetric Schur functions. Hopf algebras, quasisymmetric functions, and Young composition tableaux.* New York, NY: Springer, 2013. doi: 10.1007/978-1-4614-7300-8.

S. Mason. A decomposition of Schur functions and an analogue of the Robinson-Schensted-Knuth algorithm. *Séminaire Lotharingien de Combinatoire*, 57:B57e, 2008.

B. E. Sagan. Shifted tableaux, Schur Q-functions, and a conjecture of R. Stanley. *Journal of Combinatorial Theory, Series A*, 45(1):62–103, 1987.

B. E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, volume 203 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2001.

R. Stanley. On the number of reduced decompositions of elements of Coxeter groups. *Eur. J. Comb.*, 5:359–372, 1984.

R. P. Stanley. *Enumerative Combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999.

D. R. Worley. *A theory of shifted Young tableaux.* PhD thesis, Massachusetts Institute of Technology, 1984.